
[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
**Search:** ☒ The ACM Digital Library ☐ The Guide


## THE ACM DIGITAL LIBRARY

[Feedback](#)

master replication schema Found  
 Terms used: 129 of  
**master replication schema** 255,080

Sort results by   
 Display results

[Save](#) [Refine](#)  
[these results](#)  
[to a Binder](#)  
[with Advanced Search](#)  
☐ [Open results in a new window](#)  
[Try this search in The ACM Guide](#)

Results 1 - 20 of 129 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [next](#)

>>

### 1 [Middleware-based database replication: the gaps between theory and practice](#)



Emmanuel Cecchet, George Candea, Anastasia Ailamaki  
 June 2008 SIGMOD '08: Proceedings of the 2008 ACM SIGMOD international conference on Management of data

**Publisher:** ACM

Full text available: [Pdf](#) (353.17 KB)


Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 95, Downloads (12 Months): 224, Citation Count: 0


The need for high availability and performance in data management systems has been fueling a long running interest in database replication from both academia and industry. However, academic groups often attack replication problems in isolation, overlooking ...

**Keyw ords:** database replication, middleware, practice and experience

## 2 SIPRe: a partial database replication protocol with SI replicas

 J. E. Armendáriz-Iñigo, A. Mauch-Goya, J. R. González de Mendivil, F. D. Muñoz-Escóí  
March SAC '08: Proceedings of the 2008 ACM symposium on Applied computing  
2008

**Publisher:** ACM

Full text available:  Pdf (170.37  
KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 7, Downloads (12 Months): 32, Citation Count: 0


Database replication has been researched as a solution to overcome the problems of performance and availability of distributed systems. Full database replication, based on group communication systems, is an attempt to enhance performance that works well ...

Keyw ords: database recovery, database replication, eager database replication, generalized snapshot isolation, highly-available systems, partial database replication, scalability

## 3 Database replication with Slony-I

Ludovic Marcotte  
June 2005 Linux Journal, Volume 2005 Issue 134

**Publisher:** Specialized Systems Consultants, Inc.


Full text available:  Html (22.00  
KB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)


**Bibliometrics:** Downloads (6 Weeks): 3, Downloads (12 Months): 46, Citation Count: 0

Move up to a highly available cluster without leaving behind the open-source database you trust.

## 4 Evaluation of database management systems for Erlang

 Emil Hellman  
September ERLANG '06: Proceedings of the 2006 ACM SIGPLAN workshop on  
2006 Erlang

**Publisher:** ACM

Full text available:  Pdf (291.35  
KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 13, Downloads (12 Months): 288, Citation Count: 0

Erlang/OTP's DBMS Mnesia is lacking in several important areas to consider when implementing very large databases with massive scalability requirements. This article reveals the result from a study examining what Erlang developers consider important ...

Keyw ords: AHP, DBMS, Erlang

## 5 The ADO.NET entity framework: making the conceptual level real



José A. Blakeley, David Campbell, S. Muralidhar, Anil Nori  
December 2006 ACM SIGMOD Record, Volume 35 Issue 4

**Publisher:** ACM

Full text available: Pdf (388.24 KB)

Additional Information: [full citation](#), [abstract](#), [references](#)

**Bibliometrics:** Downloads (6 Weeks): 12, Downloads (12 Months): 111, Citation Count: 2

This paper describes the ADO.NET Entity Framework, a platform for programming against data that raises the level of abstraction from the logical (relational) level to the conceptual (entity) level, and thereby significantly reduces the impedance mismatch ...

## 6 Fedora directory server: the evolution of Linux authentication

Jeramiah Bowling  
September 2007 Linux Journal, Volume 2007 Issue 161

**Publisher:** Specialized Systems Consultants, Inc.

Full text available: Html (57.51 KB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 4, Downloads (12 Months): 2141, Citation Count: 0

Want an alternative to OpenLDAP?

## 7 Managing duplicates in a web archive



Daniel Gomes, André L. Santos, Mário J. Silva  
April 2006 SAC '06: Proceedings of the 2006 ACM symposium on Applied computing 2006

**Publisher:** ACM

Full text available: Pdf (266.54 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 7, Downloads (12 Months): 83, Citation Count: 2

Crawlers harvest the web by iteratively downloading documents referenced by URLs. It is frequent to find different URLs that refer to the same document, leading crawlers to download duplicates. Hence, web archives built through incremental crawls waste ...

## 8 Extending DBMSs with satellite databases

Christian Plattner, Gustavo Alonso, M. Tamer Özsu

July 2008 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 17 Issue 4

**Publisher:** Springer-Verlag New York, Inc.

Full text available:  Pdf (681.33 KB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 3, Downloads (12 Months): 3, Citation Count: 0

In this paper, we propose an extensible architecture for database engines where satellite databases are used to scale out and implement additional functionality for a centralized database engine. The architecture uses a middleware layer that offers consistent ...

**Keywords:** Dynamic satellite creation, Extending database functionality, Satellite databases, Snapshot isolation


## 9 The co-replication methodology and its application to structured parallel programs



Carlo Bertolli, Massimo Coppola, Corrado Zoccolo

October 2007 CompFrame '07: Proceedings of the 2007 symposium on Component and framework technology in high-performance and scientific computing

**Publisher:** ACM

Full text available:  Pdf (215.08 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 1, Downloads (12 Months): 40, Citation Count: 0

We introduce Co-Replication, a technique exploiting abstract properties of a computation to allow parallel replicas of a software module to cooperate, enhancing both the reliability and availability of the resulting component, and providing a flexible ...

**Keywords:** replication, structured parallelism

## 10 A meta model and an infrastructure for the non-transparent replication of object databases



Werner Dreyer, Klaus R. Dittrich

November 2000 CIKM '00: Proceedings of the ninth international conference on Information and knowledge management

**Publisher:** ACM

Full text available:  Pdf (179.36 KB)

Additional Information: [full citation](#), [references](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 5, Downloads (12 Months): 39, Citation Count: 0

**Keywords:** object databases, object replication, replication meta models

**11** A taxonomy of Data Grids for distributed data sharing, management, and processing



Srikumar Venugopal, Rajkumar Buyya, Kotagiri Ramamohanarao  
June ACM Computing Surveys (CSUR), Volume 38 Issue 1  
2006

**Publisher:** ACM

Full text available: Pdf (1.70 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 167, Downloads (12 Months): 1449, Citation Count: 6

Data Grids have been adopted as the next generation platform by many scientific communities that need to share, access, transport, process, and manage large data collections distributed worldwide. They combine high-end computing technologies with high-performance ...

**Keyw ords:** Grid computing, data-intensive applications, replica management, virtual organizations

**12** Model-driven development of Web applications: the AutoWeb system



Piero Fraternali, Paolo Paolini  
October ACM Transactions on Information Systems (TOIS), Volume 18 Issue  
2000 4

**Publisher:** ACM

Full text available: Pdf (6.94 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 63, Downloads (12 Months): 469, Citation Count: 22

This paper describes a methodology for the development of WWW applications and a tool environment specifically tailored for the methodology. The methodology and the development environment are based upon models and techniques already used in the hypermedia, ...

**Keyw ords:** HTML, WWW, application, development, intranet, modeling

**13** Exchanging intensional XML data



Tova Milo, Serge Abiteboul, Bernd Amann, Omar Benjelloun, Fred Dang Ngoc  
June SIGMOD '03: Proceedings of the 2003 ACM SIGMOD international conference  
2003 on Management of data

**Publisher:** ACM


Full text available: Pdf (237.21 KB)


Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 4, Downloads (12 Months): 54, Citation Count: 18

XML is becoming the universal format for data exchange between applications. Recently, the emergence of Web services as standard means of publishing and accessing data on the Web introduced a new class of XML documents, which we call *intensional* ...

#### 14 Designing data marts for data warehouses

 October 2001 ACM Transactions on Software Engineering and Methodology (TOSEM), Volume 10 Issue 4  
**Publisher:** ACM

Full text available:  Pdf (203.43 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#), [review](#)


**Bibliometrics:** Downloads (6 Weeks): 86, Downloads (12 Months): 649, Citation Count: 10

Data warehouses are databases devoted to analytical processing. They are used to support decision-making activities in most modern business settings, when complex data sets have to be studied and analyzed. The technology for analytical processing assumes ...

**Keywords:** conceptual modeling, data mart, data warehouse, design method, software quality management

#### 15 Scalable database replication through dynamic multiversioning

Kaloian Manassiev, Cristiana Amza  
October 2005 CASCON '05: Proceedings of the 2005 conference of the Centre for Advanced Studies on Collaborative research  
**Publisher:** IBM Press

Full text available:  Pdf (313.41 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 12, Downloads (12 Months): 114, Citation Count: 0

We scale the database back-end in dynamic content web servers on a set of database replicas while maintaining strong consistency. This is contrary to conventional wisdom in replicated databases which says that one could have either strong consistency ...

## 16 Ganymed: scalable replication for transactional web applications

Christian Plattner, Gustavo Alonso

October Middlew are '04: Proceedings of the 5th ACM/IFIP/USENIX international  
2004 conference on Middleware

**Publisher:** Springer-Verlag New York, Inc.

Full text available:  Pdf (295.27  
KB)


Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#)

**Bibliometrics:** Downloads (6 Weeks): 10, Downloads (12 Months): 97, Citation Count: 25

Data grids, large scale web applications generating dynamic content and database service providing pose significant scalability challenges to database engines.

Replication is the most common solution but it involves difficult trade-offs. The most difficult ...

## 17 Multiview access protocols for large-scale replication

 Xiangning Liu, Abdelsalam Helal, Weimin Du  
June ACM Transactions on Database Systems (TODS), Volume 23 Issue 2  
1998

**Publisher:** ACM

Full text available:  Pdf (365.98  
KB)


Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#),  
[index terms](#), [review](#)

**Bibliometrics:** Downloads (6 Weeks): 9, Downloads (12 Months): 64, Citation Count: 2

The article proposes a scalable protocol for replication management in large-scale replicated systems. The protocol organizes sites and data replicas into a tree-structured, hierarchical cluster architecture. The basic idea of the protocol is to accomplish ...

Keyw ords: data replication, large-scale systems, multiview access

## 18 Bigtable: A Distributed Storage System for Structured Data

 Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach, Mike Burrows, Tushar Chandra, Andrew Fikes, Robert E. Gruber  
June ACM Transactions on Computer Systems (TOCS), Volume 26 Issue 2  
2008

**Publisher:** ACM

Full text available:  Pdf (300.96  
KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 133, Downloads (12 Months): 499, Citation Count: 0

Bigtable is a distributed storage system for managing structured data that is designed to scale to a very large size: petabytes of data across thousands of commodity servers. Many projects at Google store data in Bigtable, including web indexing, Google ...

Keyw ords: Large-Scale Distributed Storage

## 19 Affinity-based management of main memory database clusters



Minwen Ji

November 2002      ACM Transactions on Internet Technology (TOIT), Volume 2 Issue 4

**Publisher:** ACM

Full text available: Pdf (553.96 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 8, Downloads (12 Months): 73, Citation Count: 1

We study management strategies for main memory database clusters that are interposed between Internet applications and back-end databases as content caches. The task of management is to allocate data across individual cache databases and to route queries ...

**Keyw ords:** Main memory database, clustering, database administration, database cluster, file organization, query affinity, scalability

## 20 The state of the art in distributed query processing



Donald Kossmann

December 2000      ACM Computing Surveys (CSUR), Volume 32 Issue 4

**Publisher:** ACM

Full text available: Pdf (455.39 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

**Bibliometrics:** Downloads (6 Weeks): 125, Downloads (12 Months): 1216, Citation Count: 45

Distributed data processing is becoming a reality. Businesses want to do it for many reasons, and they often must do it in order to stay competitive. While much of the infrastructure for distributed data processing is already there (e.g., modern network ...

**Keyw ords:** caching, client-server databases, database application systems, dissemination-based information systems, economic models for query processing, middleware, multitier architectures, query execution, query optimization, replication, wrappers

---

Results 1 - 20 of 129 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [next](#)

[>>](#)

The ACM Portal is

published by the Association for Computing Machinery. Copyright © 2008 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)